

Photowalking the City: Comparing Hypotheses About Urban Photo Trails on Flickr

Martin Becker¹(✉), Philipp Singer², Florian Lemmerich², Andreas Hotho¹, Denis Helic³, and Markus Strohmaier^{2,4}

¹ University of Wuerzburg, Wuerzburg, Germany
{becker,hotho}@informatik.uni-wuerzburg.de

² GESIS - Leibniz Institute for the Social Sciences, Mannheim, Germany
{philipp.singer,florian.lemmerich,markus.strohmaier}@gesis.org

³ Graz University of Technology, Graz, Austria
dhelic@tugraz.at

⁴ University of Koblenz-Landau, Mainz, Germany
strohmaier@uni-koblenz.de

Abstract. Understanding human movement trajectories represents an important problem that has implications for a range of societal challenges such as city planning and evolution, public transport or crime. In this paper, we focus on geo-temporal *photo* trails from four different cities (Berlin, London, Los Angeles, New York) derived from Flickr that are produced by humans when taking sequences of photos in urban areas. We apply a Bayesian approach called HypTrails to assess different explanations of how the trails are produced. Our results suggest that there are common processes underlying the photo trails observed across the studied cities. Furthermore, information extracted from social media, in the form of concepts and usage statistics from Wikipedia, allows for constructing explanations for human movement trajectories.

1 Introduction

Understanding the way people navigate urban areas represents an important problem that has implications for a range of societal challenges such as city planning and evolution, public transportation or crime. Recent research in the computational social sciences has studied human movement trajectories in cities through a variety of data sources including mobile phone data [15, 35], GPS tracking [41], WiFi tracking [29], location-based social media platforms [8], online photo sharing sites [11, 13, 14] and others. Such studies have provided a number of insights into human movement trajectories. For example, past work has indicated that human mobility exhibits regularities [15, 35] and spatio-temporal patterns [8]. Research has also shown that we can successfully leverage these patterns for certain tasks such as constructing high quality travel itineraries [11]. Yet, little is known about how the corresponding trails materialize, i.e., what factors play a role when people move through urban spaces. Thus, in this paper, we extend the stream of research on human movement by assessing different

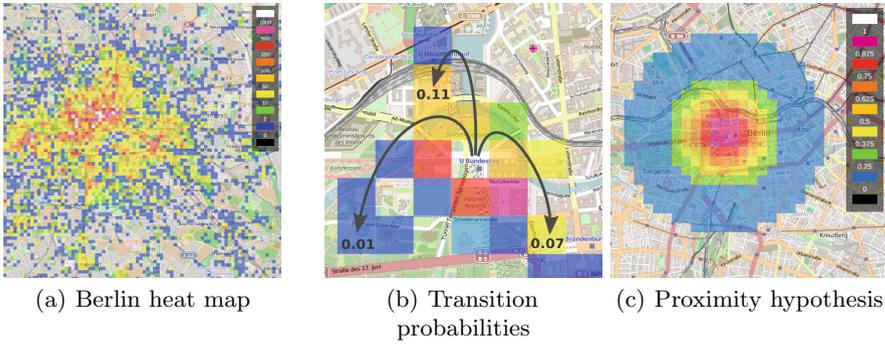


Fig. 1. *Main concepts of this work.* In (a), we visualize a cell-based grid layout of Berlin with photo frequencies visualized in a heatmap format, as derived from our data at interest. (b) depicts an example of transition probabilities between cells. In particular, it depicts the cell where the “German Bundestag” is located, and visualizes the transition probabilities to subsequent cells, i.e. cells where people take photos after they photographed the Bundestag. For instance, with a probability of 0.07, people take a picture at the “Brandenburg Gate” after they have taken one at the “Bundestag”. As we are interested in gaining insights into the processes producing these trails, we formulate hypotheses based on belief in transitions (parameters of the Markov chain). In (c), we depict an exemplary proximity hypothesis which represents a belief that people successively take photos in proximate areas of a city; higher proximity refers to higher belief.

explanations (hypotheses) of *how* urban photo trails are produced. A better understanding of this process is relevant for a series of practical problems, such as local recommendations of picturesque locations, studying touristic movement patterns or movement of people in urban environments in general.

Problem and Objective. To that end, we aim to *compare a set of different hypotheses about how urban photo trails are produced* given actual sequential photo data from four cities (Berlin, London, Los Angeles and New York) as derived from Web data (Flickr). In particular, we want to assess the plausibility of different potential explanations (hypotheses) for the trails of photos that we can observe and compare them across cities. We define an *urban photo trail* as a sequence of spatial positions in a city over a period of time as, e.g., obtained from the geo-temporal metadata of photos. Then, *hypotheses* can be expressed as different beliefs about transitions between spatial positions. For example, we might want to compare a *proximity hypothesis*—which represents a belief that people frequently take subsequent photos in geographically close regions of a city— with another *points of interest (POI) hypothesis* that represents the belief that humans take subsequent photos of POIs.

Approach. To tackle these challenges, we resort to a Bayesian approach called *HypTrails* [32] that allows for relative comparison of hypotheses given sequential categorical data. The approach is based on Markov chain modeling and Bayesian

inference; hypotheses can be seen as Markov transitions and our belief in them. For working with HypTrails, we map each photo of a trail to a cell of a geo-spatial grid over the city. For constructing hypotheses, we utilize information extracted from the social semantic web, specifically, from Wikipedia, DBpedia and YAGO and corresponding usage statistics. For illustration, we visualize and explain the main concepts of this work in Figure 1.

Contributions and Findings. The main contribution of this work is a systematic evaluation of hypotheses for explaining how urban photo trails are produced in four different cities. We find interesting commonalities, in particular we find that the partial ordering of evidence for different hypotheses is quite stable across the cities we have investigated. Furthermore, information extracted from social media, in the form of concepts and usage statistics from Wikipedia, allows for finding advanced explanations for human movement trajectories. Most prominently, our results suggest that humans seem to prefer to consecutively take photos at proximate POIs that are also popular on Wikipedia. In addition, we can also observe differences between cities: For example, proximity is less relevant for Los Angeles, which is a plausible finding given the unique topology of the city among the studied datasets.

The findings of our work can enable photo sharing websites to offer localized recommendations of picturesque photo spots according to actual tourist trajectories, city planners to explore human movement patterns of its inhabitants in general or tourist organizations to facilitate and optimize tours.

Structure. We start by giving an overview of related work in Section 2. Section 3 describes the utilized HypTrails method as well as the Flickr data at interest. In Section 4, we present our hypotheses of interest. Section 5 summarizes our experiments and presents the corresponding results. We discuss our work in Section 6 and conclude it in Section 7.

2 Related Work

This section covers related work from three areas of research: (i) studies on geo-spatial trails, (ii) studies on Flickr, and (iii) studies on human trails on the Web in general.

Geo-Spatial Trails. In the past, researchers have studied geo-spatial trails originating from diverse sources along several dimensions. For example, Song et al. [35] as well as Gonzales et al. [15] showed that geo-spatial trails (derived from mobile phone data) indicate high predictability with a lack of variability between humans. Cho et al. [8] supplemented mobile phone data with check-in data from social networks and found that human mobility is not only driven by periodicity, but also by social ties. Generally, social check-in data on the Web has supplemented our understanding of human mobility as e.g., shown in [9, 25, 26]. Further works have studied human mobility as derived from e.g., Twitter [12, 16], taxi data [27] or bike data [18].

Flickr. Geo-spatial trails have also been studied on Flickr; e.g., De Choudhury et al. [11] aimed at leveraging photo trails for automatically constructing travel itineraries through cities by utilizing popularity of POIs. Similarly, Tai et al. [36] used past landmarks photographed by users for recommending sequences of new landmarks as derived from sequential information by other users on Flickr. Girardin and colleagues have conducted several studies on Flickr photo trails. In [13], they studied digital footprints and in [14] they focused on tourist dynamics based on concentrations and spatio-temporal flows revealing popular points of interests, density points and common trails tourists follow. Apart from trails, Flickr has also been studied in other contexts specifically regarding tagging [10, 23, 31] and social network properties [5, 24].

Human Trails on the Web. Our research community has studied sequential interactions of humans with the Web in various settings. While some work has mainly focused on modeling (e.g., [2, 3, 7, 28, 30, 33]), others have been more interested in investigating regularities, patterns and strategies (e.g., [17, 38, 39]) that emerge when humans sequentially engage with the Web. This kind of work has been interested in answering similar objectives as this paper, i.e., understanding sequential steps by humans.

Many researches have been studying human navigational trails on the Web [4, 17, 33, 39]. This line of research has inspired other works to improve the Web, based on these found patterns by e.g., better website design [6]. Similarly, a better understanding of human mobility might lead to improvements in urban city design. Other works have focused on e.g., studying search trails [40], diffusion trails [1] or ontology editing trails [37].

3 Methods and Material

For comparing hypotheses about how humans choose their next location to take a photo, we resort to an approach called *HypTrails* [32] and study data derived from *Flickr*, a social photo sharing platform. This section describes both the method and the data set.

3.1 Methodology

HypTrails [32] is an approach utilizing first-order Markov chain models [33] and Bayesian inference [20] for expressing and comparing hypotheses about human trails. Since a Markov chain model is a stochastic system that models transition probabilities between a set of discrete states, it is not directly applicable to the continuous space of geo-spatial data. To address this issue, we discretize the geo-spatial regions by defining a grid over the target area where the grid cells then can be interpreted as states in the Markov chain. This step allows us to use HypTrails in the context of geo-spatial data.

With HypTrails, hypotheses can be expressed as beliefs in Markov transitions, i.e., assumptions on common and uncommon transitions at individual states; Section 4 presents several such hypotheses about human photo trails. To obtain

insights into the relative plausibility of a set of hypotheses given data, HypTrails resorts to Bayesian inference and specifically to the marginal likelihood—also called *evidence*—denoting the probability of the data D given a hypothesis H . HypTrails utilizes the sensitivity of the marginal likelihood on the conjugate Dirichlet prior for comparing hypotheses. The main idea is thus, to incorporate hypotheses as Dirichlet priors into the inference process. The hyperparameters of Dirichlet distributions can be interpreted as pseudo counts; higher pseudo counts refer to higher beliefs in a specific transition.

Thus, for each hypothesis, we need to provide a hypothesis matrix Q that captures our generic beliefs in transitions (see Section 4) based on our geo-spatial states. HypTrails then internally elicits proper Dirichlet priors from these expressed hypotheses matrices by setting the pseudo counts of the priors accordingly. An additional parameter K steers the total number of pseudo counts assigned; the higher we set it, the more we believe in a given hypothesis. Basically, this means, that with higher K , the stronger we believe in the single parameter configuration specified in Q . With lower values of K , the Dirichlet prior also assigns more probability mass to other, similar parameter configurations, thus giving the hypothesis some “tolerance”.

Based on the elicited priors for different hypotheses, HypTrails determines the marginal likelihood with respect to the empirical (observed) data. Then, we can judge the relative plausibility of two hypotheses by comparing their evidences for the same K ; higher evidences refer to higher plausibility. The fraction of the evidence of two hypotheses (priors), called Bayes factor [19], is then used for determining the strength of evidence for one hypothesis over the other. A Bayes factors can be directly interpreted as the Bayesian equivalent to a frequentist’s significance value. In this article, all Bayes factors for reported results are decisive. Therefore, we refrain from explicitly reporting them and, for simplicity, we can state that one hypothesis H_1 is more plausible compared to another hypothesis H_2 , if the evidence of H_1 is higher than the one of H_2 for the same value of K . Thus, the partial ordering based on the plausibility of respective hypotheses $\mathbf{H} = \{H_1, H_2, \dots, H_n\}$ can be determined by ranking their evidences from largest to smallest for single values of K . In this work, we report evidences on a log scale. We present corresponding results in Section 5.

Table 1. Bounding boxes and center coordinates used for data collection and hypothesis creation.

	min lon.	min lat.	max lon.	max lat.	center lon.	center lat.
Berlin	13.088400	52.338120	13.76134	52.675499	13.383333	52.516667
London	-0.5103	51.2868	0.3340	51.6923	-0.1280	51.5077
Los Angeles	-118.6682	33.7037	-118.1552	34.3368	-118.2450	34.0535
New York	-74.2589	40.4774	-73.7004	40.9176	74.0071	40.7146

3.2 Datasets

In this work, we study human phototrails through cities by analyzing data from the social photo-sharing platform Flickr¹. In this section, we first describe the data collection and its transformation into the required representation of trails and state transitions. Then, we highlight some basic characteristics of our datasets.

Data Collection. Our datasets² contain metadata—i.e., user, temporal and geo (latitude and longitude) data—about images uploaded to the Flickr platform. In particular, we focus on pictures taken in the cities of Berlin, London, Los Angeles and New York between January 2010 and December 2014. For each city, we define a bounding box, see Table 1. We acquired corresponding data by crawling Flickr’s public API. Since our analysis requires an exact position, we remove pictures with less than street-level accuracy (level 16 on the Flickr scale³).

For our analyses, we interpret the sequence of all photos of a single user as a *phototrail* ordered by the time each photo was taken, regardless of the time difference between the photos, see also Section 6. Each photo in a trail is mapped to a cell of a grid that we place over the respective city according to its geo-reference. The grid cells are then used as the discrete states required by the HypTrails approach. Since we want to capture how people move between places in a city, such as sights or train stations, we choose a cell size of 200m x 200m for our experiments. From our experience, this cell size is small enough to distinguish places close to each other and it is large enough to aggregate movement at a single place as well as to reduce the sensitivity due to GPS inaccuracies. Figure 1 shows cells of such a grid on Berlin to give an idea about the chosen granularity.

Finally, since we are only interested in the sequence of different places people photograph, we remove all self-transitions (i.e., transitions from one grid cell to itself) from the photo trails in order to account for people taking several photos at one place. Basic statistics of the processed datasets are summarized in Table 2.

Points of Interest. We work with hypotheses (see Section 4) that utilize information about points of interest (POIs) in a city. We query these POIs from

Table 2. Basic dataset statistics.

city	years	photos	cells	trails	covered cells	avg. trail length	POIs	avg. view counts
Berlin	2010-11	60,978	43,052	4,364	6,343	13.97	1,085	1,240
London	2010-14	794,535	66,444	35,101	23,694	22.64	7,228	1,272
Los Angeles	2010-14	300,373	84,014	15,357	25,834	19.56	1,462	3,654
New York	2010-14	714,549	58,065	31,246	15,232	22.87	6,002	1,511

¹ <https://flickr.com>

² Dataset access can be requested via e-mail.

³ see <https://www.flickr.com/services/api/flickr.places.findByLatLon.html>

the social semantic web, in our case DBpedia [21] and YAGO [22]. For each city, the POIs are filtered by bounding box. Also, area concepts such as “Germany” or “Berlin”, which do not correspond to actual locations (*rdf:type* equal to *yago:District108552138*), are removed.

Additionally, we quantify the importance of a POI in some hypotheses. As an approximate measure of importance we take pageview counts of the Wikipedia articles describing the POIs. For that purpose, we extracted view counts from data available at the Wikimedia download page⁴—in this work, we use the view counts for January 2012. Table 2 shows the number of POIs per city and their average view count.

4 Hypotheses

The HypTrails approach allows to compare hypotheses about human trails on discrete states—see Section 3 for details. Hypotheses are expressed by constructing matrices that reflect beliefs about transitions between such states. This section describes how several intuitions about phototrails can be expressed as hypothesis matrices Q .

For our geo-spatial setting, we define states by discretizing the continuous geo-spatial area using a grid. Then, we elicit the belief in the transition between each ordered pair of grid cells. That is, given a grid cell s_i of a user’s last photo, we specify how likely her next photo will be taken in every other cell s_j . In the HypTrails approach, higher values correspond to higher beliefs in the corresponding transitions. In this paper, we express our beliefs as local transition probabilities $P(s = s_j | s_i)$. For example, if a hypothesis assumes that a user, who took her last photo in cell s_1 , will take the next photo in cell s_2 with probability 0.5, then we set $P(s = s_2 | s_1) = 0.5$. We assign these transition probabilities between states as the values of the hypothesis matrix Q : $Q(i, j) = P(s = s_j | s_i)$. Please note that the hypothesis matrix is a stochastic matrix since each row i of Q sums to 1, i.e. $\sum_j Q(i, j) = 1$.

For simplicity, we do not directly express transition beliefs as probabilities, but rather in the form of a belief function $\bar{P}(i, j)$. This function can then be transformed into a probability distribution by multiplying it with a normalization factor $\frac{1}{Z}$ obtained by summing over all values of \bar{P} with regard to the source cell s_i :

$$P(s = s_j | s_i) = \frac{1}{Z} \bar{P}(i, j), \quad Z = \sum_{j=1}^n \bar{P}(i, j)$$

In this paper, we apply Gaussian distributions for weighting transition probabilities or factors. In this context, the elements of a hypothesis matrix Q often take very small values. For practical computational reasons, we set the value for a belief in a transition $Q(i, j) = \bar{P}(i, j)$ to 0 if the transition probability falls below a threshold. For our experiments we use 0.01 as a threshold. Furthermore, we set the beliefs in self-transitions to zero ($\bar{P}(i, i) = 0$) for all hypotheses as we

⁴ <http://dumps.wikimedia.org/other/pagecounts-raw/>

are more interested in modeling actual movement. This is in accordance to the removal of self-transitions in the observed data for our experiments. Next, we describe the hypotheses that we analyze in this paper.

Uniform Hypothesis. This hypothesis believes that each transition is equally likely assuming that users take pictures uniformly at random anywhere in the city regardless of the previous location:

$$\bar{P}_{uniform}(i, j) = 1$$

We use the uniform hypothesis as a baseline hypothesis: an informative hypothesis should at least be more plausible than the uniform hypothesis in order to express valid notions about the processes underlying human photo trails.

Center Hypothesis. Typically, the city center is the most lively part of a city and this hypothesis assumes that a user always takes her next picture near the city center regardless of the location of her last picture. To formalize this hypothesis, we use the geographic center C of the city (as listed in Table 1) and lay a two-dimensional Gaussian distribution centered at this point over the corresponding grid. Given the geographic (haversine [34]) distance $dist(C, j)$ between the city center C and the central point of the grid cell s_j , we calculate the entries of the hypotheses matrices from the following distribution:

$$\bar{P}_{center}(i, j) = e^{-\frac{1}{2\sigma^2} dist(C, j)^2}$$

We parametrize the center hypothesis with the standard deviation σ (e.g., in kilometers). A small value of σ indicates the belief that most pictures are taken very close to the city centre. If σ approaches infinity the hypothesis approximates the uniform hypothesis.

Points of Interest (POI) Hypothesis. Previous work on photo trails has shown that it is possible to automatically construct travel itineraries through a city by analyzing Flickr users behavior [11]. This suggests that humans favor points of interests—including not only tourist attractions, but also important public transportation spots or the locations of government institutions—when taking photos throughout major urban tourist areas. Thus, the POI hypothesis captures the intuition that people take a majority of pictures near such POIs. Similar to the center hypothesis, we express the attraction force of each POI with a two-dimensional Gaussian distribution. Formally, for each cell s_j and each POI $q \in Q$, we get an attraction value $G(q, j)$ that corresponds to the likelihood that q generates a picture in cell s_j according to their distance:

$$G(q, j) = e^{-\frac{1}{2\sigma^2} dist(q, j)^2}$$

Here, as before, $dist(q, j)$ describes the haversine distance between POI q and cell s_j . Then, for each cell, we aggregate the distance weighted attraction values of all POIs.

$$\bar{P}_{poi}(i, j) = \sum_{q \in Q} G(q, j)$$

In doing so, cells that contain multiple POIs have a stronger attraction to their neighboring cells. Again, we have to choose an appropriate standard deviation σ ; a small σ assumes that photos are taken directly at the point of interest, whereas a larger σ assumes that pictures are taken in the surroundings of a POI. Larger values of σ may represent the fact that people do not take pictures directly at a POI, e.g., to cover an architectural attraction fully in one picture, or they find something interesting to photograph nearby. In this work, we utilize points of interest extracted from DBpedia, see Section 3.2.

Weighted POI Hypothesis. Each city contains a large amount of potential POIs, however, not all of these are equally important. For example, the “Brandenburg Gate” is more likely to influence human trails in Berlin than the less known “Charlottenburg Gate”. We capture this intuition in a weighted POI hypothesis by approximating the importance of a POI q by the view count $views(q)$ of the Wikipedia article corresponding to this POI. If the view count of an article is very high (as e.g., for the “Brandenburg Gate”), we expect the respective POI to have a stronger influence on the sequence of image locations. We quantify this hypothesis by weighting each term of the POI hypothesis:

$$\bar{P}_{weighted_poi}(i, j) = \sum_{q \in Q} (views(q) \cdot G(q, j)).$$

Since we expect view counts to follow a power law, we also apply a sub-linear weighting scheme to avoid overemphasizing the importance of large points of interest:

$$\bar{P}_{log_weighted_poi}(i, j) = \sum_{q \in Q} (\log(views(q)) \cdot G(q, j)).$$

Proximity Hypothesis. This hypothesis—motivated by findings of previous work [8, 15, 32]—expresses the belief that the next image of a user will be taken nearby the last image. To formalize this hypothesis, we consider the haversine distances $dist(i, j)$ between the center points of two cells s_i, s_j . Then, we can again specify the respective transition probabilities by applying a two-dimensional Gaussian distribution:

$$\bar{P}_{prox}(i, j) = e^{-\frac{1}{2\sigma^2} dist(i, j)^2}$$

As before, a standard deviation σ must be specified; a small value of σ suggests a photo is more likely to be taken very close to a user’s previous photo. An example for this hypothesis is depicted in Figure 1(c) where we visualize our beliefs in transitions from one state to other states (i.e., one row of Q).

Mixture of Hypotheses. Finally, we are interested in studying the effects of a mixture of two hypotheses. Technically, we mix two hypotheses by element-wise multiplication of the corresponding hypothesis matrices. In this paper, we study two mixtures combining the intuition that people are likely to take pictures at POIs or close to the city center on one hand, but at the same time stay close to their current location for their next photo on the other hand. We can capture

this by combining the POI or center hypotheses with the proximity hypothesis. Please note that other kinds of combinations are also conceivable.

Proximate weighted POI hypothesis. First, we are combining the POI hypothesis with the proximity hypothesis, i.e., we assume that people will move to a POI to take their next photo but, instead of moving to a random POI, they choose one close by.

$$\bar{P}_{prox-(log-)weighted_poi}(i, j) = \bar{P}_{prox}(i, j) \cdot \bar{P}_{(log-)weighted_poi}(i, j)$$

Proximate center hypothesis. Similarly, the following formulation expresses the belief that the next picture is likely taken closer to the city center, but limits the area to move to to a location close to the current one.

$$\bar{P}_{prox_center}(i, j) = \bar{P}_{prox}(i, j) \cdot \bar{P}_{center}(i, j)$$

5 Experiments

In Section 4, we introduced a set of hypotheses that express beliefs on where people take their next picture while moving through a city. In this section, we compare these hypotheses with each other based on empirical trails derived from four different cities—Berlin (Germany), London (United Kingdom), Los Angeles (USA), and New York (USA) (see Section 3.2)—by employing the HypTrails approach as outlined in Section 3.1.

First, we focus on Berlin as a representative example in Section 5.1. We report in-depth experimental results for different parameter settings of each hypothesis. Afterwards, we report results for the other three cities in Section 5.2 focusing on the individually best parameter settings and highlight prominent differences between them.

5.1 Berlin

In this section, we thoroughly study each hypothesis and their different parameterizations in the same order as they have been introduced in Section 4, focusing on Berlin.

Center Hypotheses. For Berlin, the most photos are clearly centered around the cultural center as shown in Figure 1(a). Thus, we expect the center hypothesis—i.e., the belief that people move towards the city center and stay there for taking photos (see Section 4)—to be a better explanation of human photowalking behavior than our baseline (uniform) hypothesis. We consider the center of Berlin (see Table 1) and four different standard deviations σ : 1km, 3km, 5km and 10km. The results are depicted in Figure 2(a).

As expected, the results show that for all considered values of $K > 0$ and all parameterizations of the hypothesis, the center hypothesis is more plausible than the uniform hypothesis (higher evidences). The best center hypothesis is

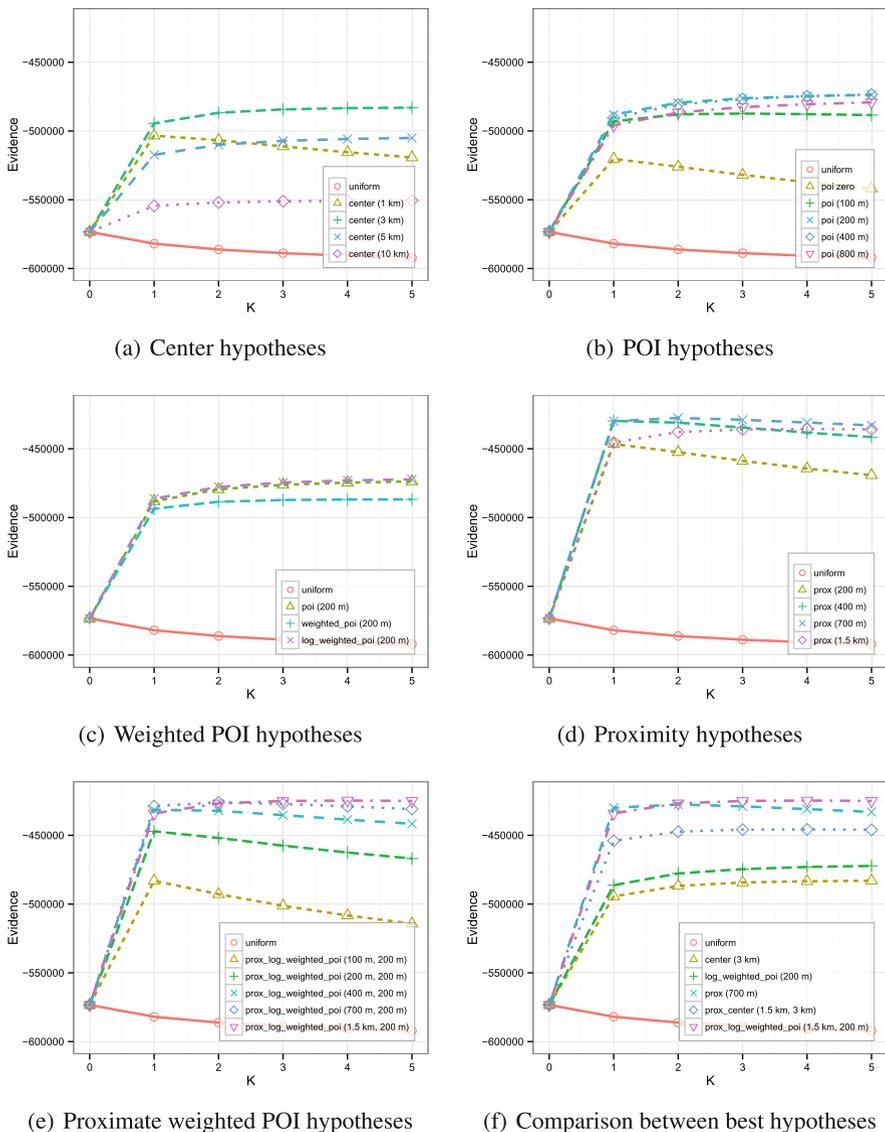


Fig. 2. Berlin. This figure visualizes the results for our hypotheses on human photo trails on Berlin. First, for each type of hypotheses at interest, we compare various parameter configurations (a-e). Then, in (f) we compare the best hypotheses from each set. Overall (f), a combination of proximity and weighted POIs provides the best hypothesis. This suggests that people prefer to subsequently take photos at important, yet proximate POIs in a city (c.f., Section 5.1).

based on $\sigma = 3\text{km}$ and the worst on $\sigma = 10\text{km}$. Standard deviations of 1km and 3km are mediocre and cross for increasing K . The initially high evidence values of 1km mean that this hypothesis covers an important aspect of the data.

The quickly dropping values, however, are an indicator that it also fails to model important transitions outside the 1km radius. This is because with increasing K , HypTrails decreases the tolerance of a hypothesis, cf. Section 3.1 and [32]. Contrary, for $\sigma = 5\text{km}$ low values of K show lower evidence, but it does not drop as quickly, eventually resulting in higher evidence values than for $\sigma = 1\text{km}$. This indicates that the 5km standard deviation covers most transitions, but fails to model the strong focus on the center aspect.

Overall, we find that the center hypothesis is a reasonable explanation for photowalking trails in Berlin. In detail, of the investigated standard deviations, 3km works best, while 1km is too specific and 5km is too broad.

Points of Interest Hypotheses. For the POI hypothesis (see Section 4), we consider five different standard deviations: 0m (only considering the grid cell the POI is located in), 100m, 200m, 400m, and 800m. The results (see Figure (2b)) suggest that the POI hypothesis provides good explanations about how people photowalk a city as all parameterizations indicate higher evidence compared to the baseline (uniform) hypothesis. In detail, the results show that the POI hypothesis without Gaussian spread performs inferior to those POI hypotheses allowing their influence to spread. The two rather close-ranged spreads 200m and 400m perform the best implying that people indeed move towards POIs. The worse performance of too narrow and too wide ranges is an indicator that people tend to visit places and take photos of the place at a close range, but not necessarily directly at the POI. For example, a minimum range might be required to capture a large building in one picture.

Weighted Points of Interest Hypotheses. In the weighted POI hypotheses, more popular POIs have a stronger influence on the transitions. For tractability, we focus on the best spreading parameter σ for the unweighted POI hypothesis from the previous paragraph, i.e. $\sigma = 200\text{m}$. Overall (see Figure (2c)), the hypothesis that people prefer to take pictures at places with many important POIs (here, derived from Wikipedia) provides a reasonable explanation for how people photowalk a city. By using online usage statistics from Wikipedia (view counts), we can strengthen the evidence of the hypothesis by a small, but significant amount if we use logarithmic scaling.

Proximity Hypotheses. For the the proximity hypothesis (see Section 4), we use four different standard deviations σ : 200m, 400m, 700m and 1.5km. Overall, the results shown in Figure (2d) demonstrate that the hypothesis that people prefer to consecutively take pictures in their proximity captures an important aspect of the production of human photo trails; $\sigma = 700\text{m}$ produces the highest evidence for all considered values of $K > 0$. For standard deviations of 200 and 400m, a similar situation occurs as for the center hypotheses with a standard deviation of 1km: They seem to concentrate their belief on a too narrow proximity leading to decreasing evidence values for higher values of K . Contrary, the proximity hypothesis with $\sigma = 1.5\text{km}$ is too broad, somewhat neglecting the centralized character of the proximity aspect.

Mixtures of Hypotheses. To evaluate the mixture of POI and proximity hypothesis, we focus on the logarithmically weighted POI hypothesis with $\sigma = 400\text{m}$ since it was one of the best performing hypotheses so far. This is combined with different standard deviations for proximity, i.e., 100m, 200m, 400m, 700m and 1.5km. The results shown in Figure (2e) indeed demonstrate that adding the proximity aspect to the POI hypothesis strongly improves the evidence of the corresponding belief how people consecutively take pictures in a city. The best results can be achieved with larger standard deviations σ , i.e., $\sigma = 700\text{m}$ and $\sigma = 1.5\text{km}$.

We also investigated different parametrization for the mixture of the proximity and the center hypothesis (not visualized in this paper due to limited space). The best parameter setting was a standard deviation of 3km for the center and a standard deviation of 1.5km for the proximity hypothesis (see Figure (2f)).

Comparison. For a direct comparison of the different hypotheses we are taking the most plausible ones (best parameters) of each set as elaborated beforehand. The results are shown in Figure (2f). We can see that the center and the *weighted* POI hypothesis perform quite similar which may be due to the larger number of (important) POIs in the city center. At the same time, the proximity hypothesis performs very well and combining it with the other hypotheses improves them strongly. Indeed, the combination of the proximity hypothesis and the *weighted* POI hypothesis provides the best explanation of how people move around Berlin while taking photos. This result suggests that information extracted from the social semantic web, in the form of concepts and usage statistics from Wikipedia, allows for finding advanced explanations for human movement trajectories.

5.2 Los Angeles, London and New York

To further augment the results from Section 5.1, we analyze three more cities, namely, Los Angeles (USA), London (United Kingdom) and New York City (USA). We show similarities and highlight some differences between the cities. For a concise presentation, we focus on the best parameter settings for each hypothesis only. The best parameters were determined separately for each city. Results are depicted in Figure 3. For most parts, all hypotheses perform very similar and the best parametrizations are consistent. This indicates that the hypotheses about photo trails in Berlin can be generalized to other cities quite well, implying that some basic patterns exist that even hold across countries.

However, there are two exceptions which are worth mentioning. First, in Los Angeles (see Figure (3c)), the most plausible center hypothesis has a standard deviation of 10km instead of 3km. This indicates that LA either has a very large center or none at all—arguably, LA is a spread out city which may cause this divergence. Additionally, in LA higher standard deviations for the POI hypothesis, i.e., 400m instead of 200m, are favored compared to the other cities. Also, even the best performing hypotheses are strongly decreasing with increasing K . This further supports the idea that LA is structurally different from the other cities. Second, the linearly weighted POI hypothesis in London is superior to the

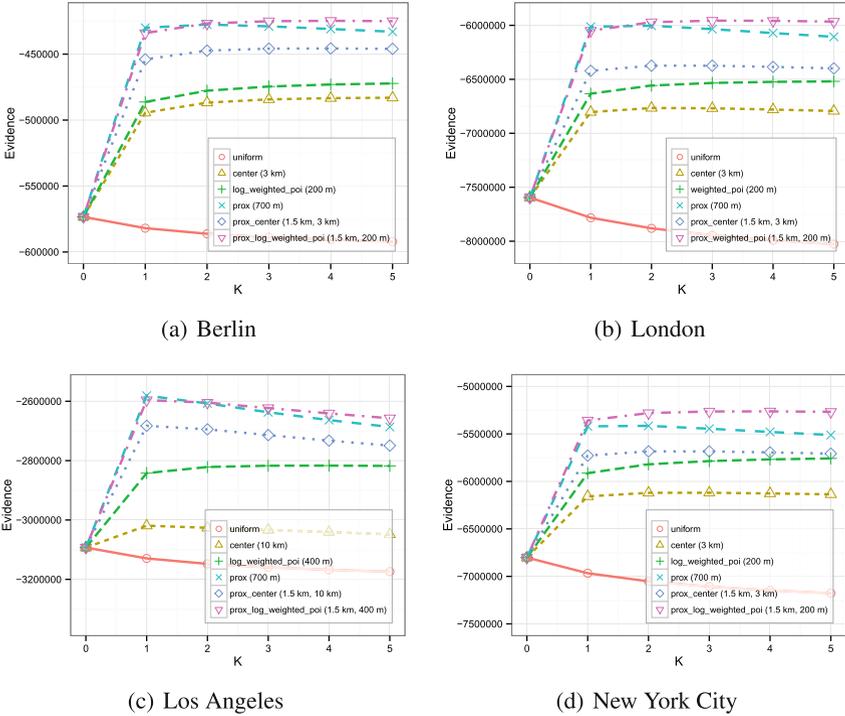


Fig. 3. *Comparison of cities.* This figure visualizes the results for our studies on human photo trails derived from Berlin(a), London (b), Los Angeles (c), and New York City (d). We only present a comparison of the best hypotheses for each type of hypotheses for each city. We can identify similar explanations across cities, but there are some differences for LA and London (c.f., Section 5.2).

logarithmically weighted one. This may be due to different view count distributions and has to be further investigated in the future.

6 Discussion

In this work, we have conducted extensive experiments to gain a better understanding of the underlying processes that are employed when people take photos while moving through cities. We dedicate this section to discuss characteristics specific to our approach and corresponding results and to highlight some potential limitations.

Data Characteristics. Next, we shortly discuss four relevant aspects regarding our data: (i) splitting photo trails due to time constraints, (ii) observation sparsity, (iii) Flickr movement characteristics and (iv) state granularity.

(i) We have considered the sequence of *all* photos of a user as a single photo-trail regardless of the time span in between two photos. However, if such a time span is too long (e.g., a week or even a few hours), the corresponding two photos are most likely unrelated. Thus, in additional experiments, we have removed

transitions with time intervals exceeding 6 or 24 hours respectively. The results are very similar to the ones reported in Section 5 which is why we refrain from explicitly reporting them.

(ii) Since we are using grids with 200m by 200m cells over relatively large areas, the number of observations for corresponding transitions is limited. However, as HypTrails automatically focuses on observed states, the sparsity of the data does not randomly bias our results. Thus, no further testing of possibly derivable predictors is necessary since all available information is drawn directly from the data via Bayesian inference.

(iii) Due to our focus on studying Flickr, we are only able to make judgments about behavioral aspects that emerge when people move through a city and take photos as captured by Flickr. Studying other forms of mobility data might reveal different results. However, we might assume that certain behavioral aspects are similar, regardless which type of data we look at as suggested in [8]. This may be verified, for example by contrasting different cities or consider different kinds of movement data—e.g., social check-in data, mobile phone data or business reviewing data.

(iv) While the focus on an intra-city level in this work has allowed us to gain insights into urban behavior, we might observe different movement patterns if we extended our scope of interest. To give an example, by looking at cities, we constrain our studies to a small geographic area which might favor proximity based hypotheses. If we extended the scope, by e.g., looking at a country or continent level, the results might largely differ. However, then, other kinds of hypotheses might be more plausible to study.

Choice of Hypotheses. The observations in this work are limited by our choice of which hypotheses to study and how to express them; they have mostly been motivated by related work. Many other kinds of hypotheses are conceivable and can be investigated with HypTrails and our data. We suggest some potential candidates: (i) A hypothesis expressing the belief that a river is a natural barrier in a city. (ii) Also, district boundaries may be some kind of barrier. Additionally, (iii) demographic aspects (such as crime rates) might influence movement patterns in a city.

Tourists and Other User Groups. Previous work has suggested that the photographing behavior on Flickr differs between tourists and residents in a city [11,14]. The authors of [11] argue that residents are not under the direct pressure of visiting as many POIs within a certain time span as tourists are. Thus, we might also see differences in their behavioral aspects producing the human photo trails studies in this article. The same may apply for a number of other user groups or sub-groups, such as visitors from different countries or users from different generations. While we have focused on an aggregated view in this paper, the distinction between such user groups might be an additional highly interesting layer which we leave open for future work.

7 Conclusion

In this paper, we have investigated and compared hypotheses about urban photo trails across different cities by analyzing sequences of geotagged photos uploaded to the Flickr platform using the Bayesian HypTrails approach. For the informed specification of hypotheses, we have utilized additional data sources such as DBpedia, YAGO and view counts of Wikipedia articles which has allowed us to find advanced explanations for human movement trajectories. Our results suggest that cities share interesting commonalities and differences. For example, while proximity is an overall good explanation across all cities, for the city of Los Angeles we observe movement patterns on a different scale. Most prominently, our results suggest that humans seem to prefer to consecutively take photos at proximate POIs that are popular on Wikipedia.

In future work, we plan on extending our work by looking at additional cities. Furthermore, it would be interesting to expand the current city-level analysis to a larger scale, e.g., trails across different cities or countries. Finally, improved tool support for the interactive exploration of location sequences and hypotheses would be helpful.

Acknowledgments. This work was partially funded by the DFG in the research project “PoSTs II”. Special thanks goes to Georg Dietrich from Chair 6 of Computer Science at the University of Würzburg for crawling the Flickr data.

References

1. An, J., Quercia, D., Crowcroft, J.: Partisan sharing: facebook evidence and societal consequences. In: Conference on Online Social Networks, pp. 13–24. ACM (2014)
2. Borges, J., Levene, M.: Evaluating variable-length markov chain models for analysis of user web navigation sessions. *IEEE Transactions on Knowledge and Data Engineering* **19**(4), 441–452 (2007). <http://dx.doi.org/10.1109/TKDE.2007.1012>
3. Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. In: International Conference on World Wide Web, pp. 107–117. Elsevier Science Publishers B. V. (1998)
4. Catledge, L.D., Pitkow, J.E.: Characterizing browsing strategies in the world-wide web. *Computer Networks and ISDN Systems* **27**(6), 1065–1073 (1995)
5. Cha, M., Mislove, A., Gummadi, K.P.: A measurement-driven analysis of information propagation in the flickr social network. In: International Conference on World Wide Web, pp. 721–730. ACM (2009)
6. Chi, E.H., Pirolli, P., Pitkow, J.: The scent of a site: a system for analyzing and predicting information scent, usage, and usability of a web site. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 161–168. ACM (2000)
7. Chi, E.H., Pirolli, P.L.T., Chen, K., Pitkow, J.: Using information scent to model user information needs and actions and the web. In: Conference on Human Factors in Computing Systems, pp. 490–497. ACM (2001). <http://doi.acm.org/10.1145/365024.365325>

8. Cho, E., Myers, S.A., Leskovec, J.: Friendship and mobility: user movement in location-based social networks. In: International Conference on Knowledge Discovery and Data Mining, pp. 1082–1090. ACM (2011)
9. Cramer, H., Rost, M., Holmquist, L.E.: Performing a check-in: emerging practices, norms and 'conflicts' in location-sharing using foursquare. In: Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services, pp. 57–66. ACM (2011)
10. Crandall, D.J., Backstrom, L., Huttenlocher, D., Kleinberg, J.: Mapping the world's photos. In: International Conference on World Wide Web, pp. 761–770. ACM (2009)
11. De Choudhury, M., Feldman, M., Amer-Yahia, S., Golbandi, N., Lempel, R., Yu, C.: Automatic construction of travel itineraries using social breadcrumbs. In: Conference on Hypertext and Hypermedia, pp. 35–44. ACM (2010)
12. Gallegos, L., Lerman, K., Huang, A., Garcia, D.: Geography of emotion: Where in a city are people happier? arXiv preprint [arXiv:1507.07632](https://arxiv.org/abs/1507.07632) (2015)
13. Girardin, F., Calabrese, F., Fiore, F.D., Ratti, C., Blat, J.: Digital footprinting: Uncovering tourists with user-generated content. *Pervasive Computing* **7**(4), 36–43 (2008)
14. Girardin, F., Fiore, F.D., Ratti, C., Blat, J.: Leveraging explicitly disclosed location information to understand tourist dynamics: a case study. *Journal of Location Based Services* **2**(1), 41–56 (2008)
15. Gonzalez, M.C., Hidalgo, C.A., Barabasi, A.L.: Understanding individual human mobility patterns. *Nature* **453**(7196), 779–782 (2008)
16. Hawelka, B., Sitko, I., Beinat, E., Sobolevsky, S., Kazakopoulos, P., Ratti, C.: Geo-located twitter as proxy for global mobility patterns. *Cartography and Geographic Information Science* **41**(3), 260–271 (2014)
17. Huberman, B.A., Pirollo, P.L.T., Pitkow, J.E., Lukose, R.M.: Strong regularities in world wide web surfing. *Science* **280**(5360), 95–97 (1998). <http://www.sciencemag.org/content/280/5360/95.abstract>
18. Kaltenbrunner, A., Meza, R., Grivolla, J., Codina, J., Banchs, R.: Bicycle cycles and mobility patterns-exploring and characterizing data from a community bicycle program. arXiv preprint [arXiv:0810.4187](https://arxiv.org/abs/0810.4187) (2008)
19. Kass, R.E., Raftery, A.E.: Bayes factors. *Journal of the American Statistical Association* **90**(430), 773–795 (1995)
20. Kruschke, J.: *Doing Bayesian data analysis: A tutorial introduction with R*. Academic Press (2010)
21. Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P.N., Hellmann, S., Morsey, M., van Kleef, P., Auer, S., et al.: Dbpedia-a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web* (2014)
22. Mahdisoltani, F., Biega, J., Suchanek, F.: Yago3: a knowledge base from multilingual wikipedias. In: 7th Biennial Conference on Innovative Data Systems Research, CIDR 2015 (2014)
23. Marlow, C., Naaman, M., Boyd, D., Davis, M.: Ht06, tagging paper, taxonomy, flickr, academic article, to read. In: Conference on Hypertext and Hypermedia, pp. 31–40. ACM (2006)
24. Mislove, A., Koppula, H.S., Gummadi, K.P., Druschel, P., Bhattacharjee, B.: Growth of the flickr social network. In: Workshop on Online Social Networks, pp. 25–30. ACM (2008)
25. Noulas, A., Scellato, S., Mascolo, C., Pontil, M.: An empirical study of geographic user activity patterns in foursquare. *ICWSM* **11**, 70–573 (2011)

26. Noulas, A., Shaw, B., Lambiotte, R., Mascolo, C.: Topological properties and temporal dynamics of place networks in urban environments. [arXiv:1502.07979](https://arxiv.org/abs/1502.07979) [cs.SI] (2015)
27. Peng, C., Jin, X., Wong, K.C., Shi, M., Liò, P.: Collective human mobility pattern from taxi trips in urban area. *PloS one* **7**(4), e34487 (2012)
28. Pirolli, P.L.T., Pitkow, J.E.: Distributions of surfers' paths through the world wide web: Empirical characterizations. *World Wide Web* **2**(1–2), 29–45 (1999)
29. Rekimoto, J., Miyaki, T., Ishizawa, T.: LifeTag: WiFi-based continuous location logging for life pattern analysis. In: Hightower, J., Schiele, B., Strang, T. (eds.) *LoCA 2007*. LNCS, vol. 4718, pp. 35–49. Springer, Heidelberg (2007)
30. Sen, R., Hansen, M.: Predicting a web user's next access based on log data. *Journal of Computational Graphics and Statistics* **12**(1), 143–155 (2003). <http://citeseer.ist.psu.edu/sen03predicting.html>
31. Sigurbjörnsson, B., Van Zwol, R.: Flickr tag recommendation based on collective knowledge. In: *International Conference on World Wide Web*, pp. 327–336. ACM (2008)
32. Singer, P., Helic, D., Hotho, A., Strohmaier, M.: Hyptrails: a bayesian approach for comparing hypotheses about human trails on the web. In: *International Conference on World Wide Web* (2015)
33. Singer, P., Helic, D., Taraghi, B., Strohmaier, M.: Detecting memory and structure in human navigation patterns using markov chain models of varying order. *PloS one* **9**(7), e102070 (2014)
34. Sinnott, R.W.: Virtues of the haversine. *Sky and Telescope* **68**(2), 158 (1984)
35. Song, C., Qu, Z., Blumm, N., Barabási, A.L.: Limits of predictability in human mobility. *Science* **327**(5968), 1018–1021 (2010). <http://www.sciencemag.org/cgi/content/abstract/327/5968/1018>
36. Tai, C.H., Yang, D.N., Lin, L.T., Chen, M.S.: Recommending personalized scenic itinerary with geo-tagged photos. In: *International Conference on Multimedia and Expo*, pp. 1209–1212. IEEE (2008)
37. Walk, S., Singer, P., Noboa, L.E., Tudorache, T., Musen, M.A., Strohmaier, M.: Understanding how users edit ontologies: comparing hypotheses about four real-world projects. In: *International Semantic Web Conference* (2015)
38. Walk, S., Singer, P., Strohmaier, M.: Sequential action patterns in collaborative ontology-engineering projects: a case-study in the biomedical domain. In: *International Conference on Conference on Information & Knowledge Management*. ACM (2014)
39. West, R., Leskovec, J.: Human wayfinding in information networks. In: *International Conference on World Wide Web*, pp. 619–628. ACM (2012). <http://doi.acm.org/10.1145/2187836.2187920>
40. White, R.W., Huang, J.: Assessing the scenic route: measuring the value of search trails in web logs. In: *Conference on Research and Development in Information Retrieval*, pp. 587–594. ACM (2010)
41. Wiehe, S.E., Carroll, A.E., Liu, G.C., Haberkorn, K.L., Hoch, S.C., Wilson, J.S., Fortenberry, J.D.: Using gps-enabled cell phones to track the travel patterns of adolescents. *International Journal of Health Geographics* **7**(1), 22 (2008)